What we got wrong with the Internet…and how to right it.

Jon Crowcroft

Jon.crowcroft@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Network Architecture is Hard

Parable of the Zen Cyclist

Once, there was a Zen Master in a small, eldritch, fen-land university town, studied cycling for many years. One day he was ready to ride…however by This time, he was advanced in years, and had gone blind. His acolytes begged him not to go but to their amazement, he rode safely from the Castle to the Mill without even any close collisions. He continued to do this for many years until one day …

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Zen and the Art of Network Design

…he was killed in a terrible accident with another Blind Zen Master Cyclist…

Network Architecture and Design is a bit like this

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

If we don't learn from History…

- We are doomed to redesign the Model T Ford…
- The Internet has its origins in the thinking of 3 groups of people
  - Paul Baran at Rand and Donald Davies at NPL
  - Vint Cerf, Bob Kahn and Peter Kirstein
  - Ted Nelson and Tim Berners-Lee at CERN
- Between them, they devised
  - Packet switching and Fate Sharing
  - Separation of Computer and Network Concerns
  - The integration of the network into the desktop

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Fate Sharing

- The original design goal of the Internet was to split the components of a network to minimise dependencies – this led to the End-to-End principle, also known as fate sharing, due to Clark, Saltzer and Reed at MIT:
- There should be a minimum of information shared between end-system computers, and the network components that interconnect them (and amongst the network components themselves too) to promote a maximum level of survivability – i.e. if you blow up 50% of the net, communications should still be possible.

Names, Addresses and Routes #1

- We use names to distinguish what things are – the Internet Domain Name System does this too.
- The Name space is organisationally hierarchical, and has a root (1), which delegates authority downwards.
- The Name system is extensible, but not particularly dynamic (at least it wasn't designed to be so!)
- Thus www.cl.cam.ac.uk ought to correspond to roughly the same ***thing*** most of the time
- Names are stored in Name Servers, which are therefore an important operational componet…they help do this:
- Names must be mapped to addresses …

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Names, Addresses and Routes #2

- Addresses tell you where something is.
- Network addresses need to be efficiently encoded and it is common to use hierarchical addresses rooted in some authority which delegates to more local authorities
  - Topological
  - Provider
  - Geographical
- We don't necessarily need to know exactly ***where*** an address is, but we do need to be able to find out in reasonable time!
- Of course, if we move, we need to change our address or leave some kind of forwarding information
- Addresses then have to be mapped to routes …

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Names, Addresses and Routes #3

- Routes are instructions for how to get somewhere from somewhere else (usually “here”, but not always!).
- In the Internet, at any given *time*, there is typically only one route from somewhere (source address) to somewhere else (destination address). End systems give the routers these addresses as a parameter to their communication (i.e. in packets!)
- Routers are specialised computers that maintain up-to-date information about the set of routes to everywhere. To do this, they run protocols and algorithms , which can be quite complex.
- Routers are therefore a critical component.

What can go wrong?

- Links and routers in the net can break for a short or long time. If short, end systems provide recovery. If long, new route needs to be found (and end systems provide recovery).
- Protocols can be mis-implemented:- Postel's principle:
- "Be conservative in what you send, and liberal in what you receive"
- Dave Katz at Cisco says it takes 3 months to implement a router, and 3-30 years to debug it against other peoples' broken ones.
- Sometimes, this can involve subtle things like randomization of timers to avoid synchronisation effects…

TCP/IP/UDP/IP

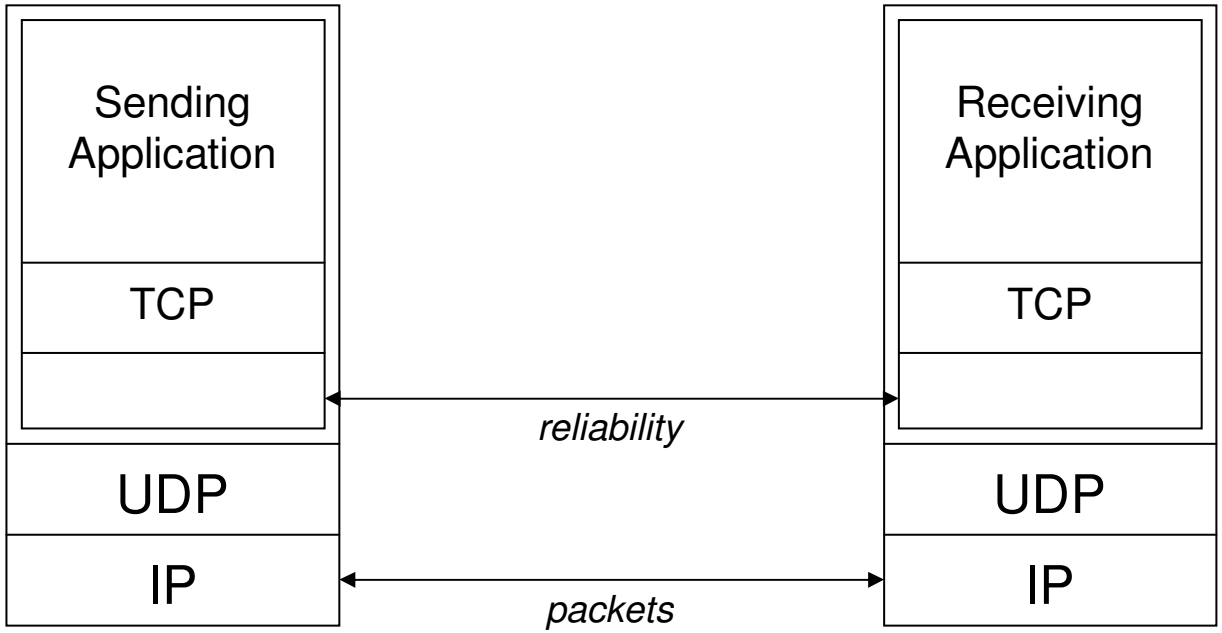
Sending Application
TCP
UDP
IP
Receiving Application
TCP
UDP
IP
reliability
packets

What is a network worth?

- The Internet is a network of networks, once called the catenet, for “concatenation”.
- Each network is part of an Autonomous System (itself possibly a network of networks) run by an Internet Service Provider.
- Metcalfe claims the value of the net is the square of number of systems.
- The topology of the network doesn’t quite reflect this (perhaps it does) – the degree of vertices varies as a power law i.e. it is roughly a Small World graph.
- suggests this is to do with economics of connection:
- New users/networks/ISPs go for connection to the rest of the net at the best connected point, that is then near to them, rather than the nearest point and then upgrade it)

But what does it cost?

- Most ISPs don't charge by volume.
- So it costs nothing to send.
- But it does cost to receive (viz cell phones in the US and the principle of unintended consequences – users left their phones off most the time)
- Leads to 2 problems:
  - Distributed Denial of Service attacks
  - SPAM
- DDoS is a network level problem, SPAM is an Application Level Problem, both causing immense cost to end user

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

DDoS and SPAM Prevention #1

- A lot of traffic *with the roughly same content* from one source to many destinations v. from many sources to one destination – what could we do?
- Maybe it is legitimate – viz Flash Crowd, Slashdot, etc -> Many-to-one
- Maybe it is fine – viz advertising/broadcast etc
- First, we need to check if the source is “who they say they are”
- Then the receiver (receiving network or end host, or human user) needs tools in their hand to prevent overload

DDoS and SPAM Prevention #2

- Need *reverse* mappings (route to address, address to name) to work correctly if we are to check things
  - First requires ISPs to check, and provide trace-back.
  - Second is part of DNS, but has to work – recently, Verisign messed this up by changing root to respond to all missing name lookups with a web page advertising their services – broke part of SPAM Prevention, and led to invoking another principle …

Excess Expressiveness #1

- Much of the Internet Architecture has been overly extensible – the DNS is one great place for playing –
- After the Verisign debacle, it has become common to quote the Principle of Least Astonishment
- “When we add some feature to a component of the Internet, it might be nice if it doesn’t impact in weird and wonderful ways on the normal operations of any and every other feature”

Excess Expressiveness #2

- A variety of other sub-systems (protocols, component services, etc) of the Internet have seen hyperactive behaviour
- HTTP, XML, SIP, BGP, etc etc
- These may just be acronyms to some people, but the working of the net depends on them
- To some extend, deployment barriers (inertial, largely) prevent complete randomness

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Economics and Choice

- Current architecture does not really allow end-user explicit choice
- Some do use multi-homing
- Its not for the bloke on the top of the Clapham Omnibus
- We'd like to switch provider to choose route (viz long-haul telephone provider choice, or TV content and provider choice, even if bundled)
- In fact, current multi-homing Is breaking routing!
- Another example of principle of unintended consequences: would be nice to fix this properly

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

What are we doing about it? #1

- Have a number of projects with some ideas on this
- Better routing and addressing (especially dealing with expected massive growth in end systems over next 10 years, with ubiquitous, sentient, ambient, pervasive computing, mobile PDAs/sensors etc etc etc)
- Invoke more (some new) science:
- Control theory, graph theory, information theory, even field theory! To try to tackle problems at the next level of scale
- For example, lets look at DDoS and SPAM:

What are we doing about it? #2

- Problem is cost to receiver – how to turn this around?
- Could use law or money – heavy handed:-
- What about computation?
  - Challenge unsolicited sender (automatically) to
  1. provide credentials (existence even!)
  2. Prove they care (e.g. do some work – for example, solve a modest 10 second cryptographic problem specific to this task)
- Apply this recursively to end-to-end, hop-by-hop so:
- Routers do it to each other, ISPs to each other, and users
- After credentials proven, add user to an “approved list” for some period
- Retains model of low number of RTTs before communications established ,which is one of the critical things the Internet got right!
- (i.e. don’t just revert to “connection oriented networking”

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

What are we doing about it? #3

- More generally, we are looking at the partitioning of functionality that is implicit in the Internet Architectural model (baroque, gothic, perpendicular?☺
- Plutarch is a framework for a more open extensible architecture – right now, we are working on naming and routing in a more contextual manner
- Xenoservers and Futuregrid are two higher level projects looking at the public computing platform space and trying to tack resource management…

Future

- Frank Zappa, the well known American composer once said “Scientists claim that hydrogen is the most abundant substance in the Universe – they are wrong – Human Stupidity is far more abundant”
- Tim Harris in the Computer Lab has suggested that our search should therefore be for the fundamental binding particle of intelligence –
- The “Clue-on”

All done!

- Thanks for listening!
- Questions?

UNIVERSITY OF CAMBRIDGE
Computer Laboratory